A Closer Look at SSD Data Integrity Requirements

Andy Tomlin

VP Firmware & Software

Problem Statement

Data Integrity in an SSD has similarities and differences with other storage systems

Error Correction (ECC)	Similar to HDDs, although differing defect types may result in different preferred correction codes
End-to-end Protection (CRC)	Typically only used on Enterprise HDDs & SSDs
Correct Address Translation (LBA)	Not a big problem on disk drives, major challenge on SSD's. Solutions may be coupled with end to end protection
Correct Version of Data	Old vs. new data selection from block recycling only performed in SSDs

Protection Levels Inside SSD

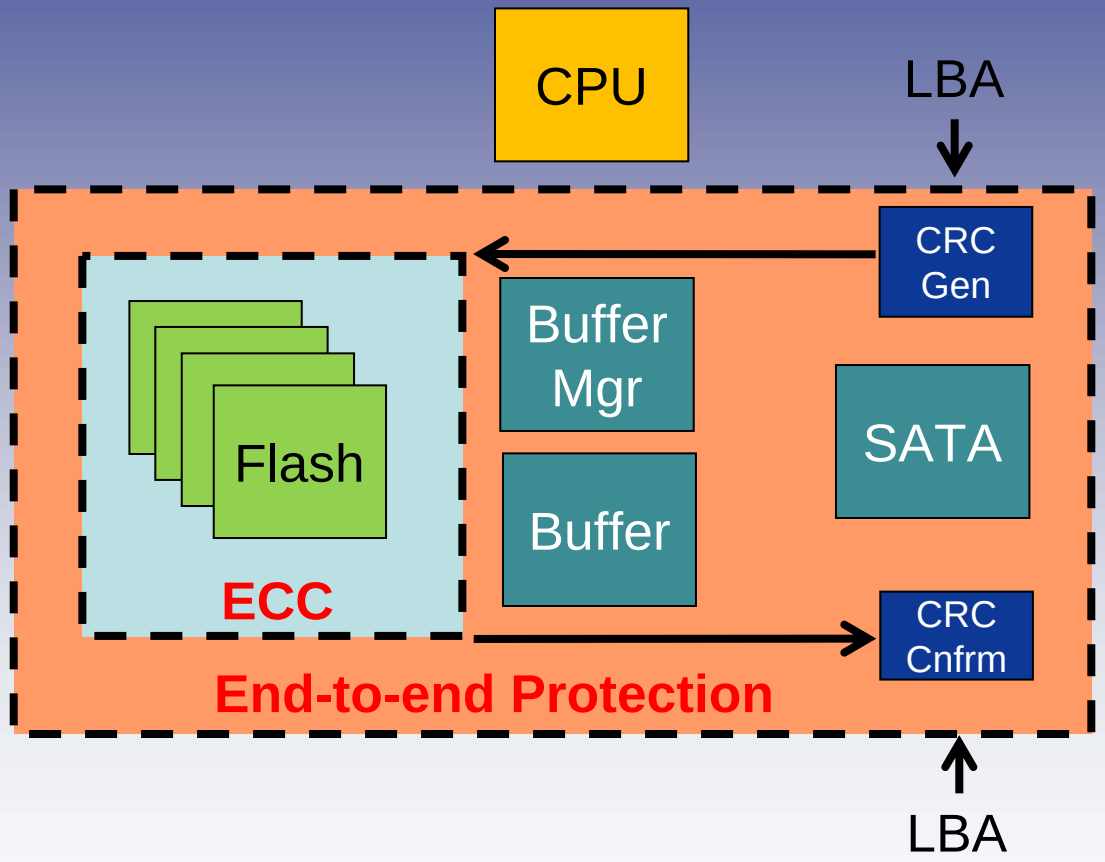

Error Correction

- All SSDs have some level of ECC on the Flash
- Typically BCH or Reed Solomon
- Requirements vary depending on
  - Process technology (5x, 4x, 3x nm)
  - Bits per cell – SLC, MLC (D2, D3, D4)
- ECC protects from Read disturb, Program Disturb, and endurance and retention page level effects
- Does not protect from Block level failure
  - More advanced controller design required for block level failures, typically required for Enterprise storage

Higher UBER Required for SSDs

- Uncorrectable Bit Error Rate must be higher for SSDs due to higher transfer rates
- High-end HDDs in the enterprise provide <1 sector error per 10^16 bits read
  - Sufficient for HDDs at 50-100 MB/s transfer rate
- SSDs that transfer 250 MB/s would show up to 5x the errors with this UBER
- SSDs will require protection to 10^17 bits read

End-to-End Protection

- Flash errors are not the only source of data integrity issues
- Modern controllers have large RAMs requiring ECC detection / correction
- Hardware and Firmware bugs can result in incorrect transfer of data from flash or address translation errors
- These types of errors are undetectable without some form of end to end protection, typically some form of CRC seeded with LBA

Correct Version of Data

- All SSDs can suffer from a problem of returning old data
- Assumes address translation functions correctly, returning old version of correct LBA. Not detectable with end to end solutions.
- Typically induced by power failure coupled with Firmware bug
- No simple solution
  - Super Cap solutions may be effective for Enterprise systems
  - Requires extensive, directed testing

Testing Methodologies For Validation of Correct Version of Data

- Must detect address translation and address versioning issues

- Data tagging
  - LBA to detect address translation errors
  - LBA versioning
    - Incrementing count for every command
    - Test system maintains table of count indexed by LBA. Note that this can create test infrastructure challenge: 512G SSD with 2Byte count per LBA = 1G RAM

Summary

- Data Integrity in an SSD has similarities and differences with other storage systems
- All SSDs have ECC protection, but new Flash generations will require higher levels of protection
- Enterprise SSDs will require UBER of 10^17 due to high transfer rates
- CRC and LBA checking can provide end-to-end protection for enterprise environments
- Validation of “correct version of data” can only be done with directed testing